

Higher-order asymptotics for the parametric complexity

James G. Dowty

October 30, 2015

Abstract

The parametric complexity is the key quantity in the minimum description length (MDL) approach to statistical model selection. Rissanen and others have shown that the parametric complexity of a statistical model approaches a simple function of the Fisher information volume of the model as the sample size n goes to infinity. This paper derives higher-order asymptotic expansions for the parametric complexity, in the case of exponential families and independent and identically distributed data. These higher-order approximations are calculated for some examples and are shown to have better finite-sample behaviour than Rissanen's approximation. The higher-order terms are given as expressions involving cumulants (or, more naturally, the Amari-Chentsov tensors), and these terms are likely to be interesting in themselves since they arise naturally from the general information-theoretic principles underpinning MDL. The derivation given here specializes to an alternative and arguably simpler proof of Rissanen's result (for the case considered here), proving for the first time that his approximation is $O(n^{-1})$.

1 Introduction

The minimum description length (MDL) principle provides a general information-theoretic approach to model selection and other forms of statistical inference [5, 17]. The MDL criterion for model selection is consistent, meaning that it will select the data-generating model from a countable set of competing parametric models with probability approaching 1 as the sample size n goes to infinity [4]. For example, if each of the parametric models is a logistic regression model with predictor variables taken from a fixed set of potential predictors, then the MDL model-selection criterion will choose the correct combination of predictors with probability approaching 1 as $n \rightarrow \infty$. The MDL model-selection criterion also has a number of strong optimality properties, which greatly extend Shannon's noiseless coding theorem [5, §III.E].

In its simplest form, the MDL principle advocates choosing the model for which the observed data has the shortest message length under a particular prefix code defined by a minimax condition [11, §2.4.3]. Shtarkov [19] showed that this is equivalent to choosing the model with the largest normalized maximum likelihood (NML) for the observed data. Here, minus the logarithm of the NML for a model \mathcal{M} and n observations x^n is

$$-\log p_n(x^n, \hat{\theta}(x^n)) + \text{Comp}$$

where: $p_n(x^n, \theta)$ is the likelihood function for data x^n at model parameter θ ; $\hat{\theta}(x^n)$ is the maximum likelihood estimate of θ corresponding to x^n ; and the parametric complexity of

the model \mathcal{M} is

$$\text{Comp} = \log \int p_n(x^n, \hat{\theta}(x^n)) d\mu_n(x^n), \quad (1)$$

where the integral is over all possible values of the data x^n (and is technically the Lebesgue integral with respect to some measure μ_n , so the integral becomes a sum if μ_n is discrete).

The NML balances how well the model fits the data (quantified by the maximized likelihood) against how complex or versatile the model is (quantified by the parametric complexity), so the NML gives a natural measure of parsimony. However, in many cases of interest, it is not practical to calculate the parametric complexity directly from the definition (1). For example, the data space for a logistic regression model consists of all binary sequences x^n of length n , and the parametric complexity is defined in terms of the sum of $p_n(x^n, \hat{\theta}(x^n))$ over this space, so calculating this sum directly is infeasible even for $n = 100$ since there are $2^{100} \approx 1.27 \times 10^{30}$ terms. Clarke and Barron [7, 8] therefore obtained $o(1)$ approximations to the parametric complexity in the limit $n \rightarrow \infty$ for independent and identically distributed (IID) data, and Rissanen [15] greatly extended these results to the non-IID setting.

Approximations to Comp are only sensible when Comp is finite, but this condition fails in many cases. In proving their $o(1)$ approximation, Clarke and Barron [7, 8] and Rissanen [15] therefore effectively restricted the data to those data points x^n whose corresponding maximum likelihood estimates $\hat{\theta}(x^n)$ lie in a given compact (i.e., closed and bounded) subset K of the parameter space Θ . Restricting x^n to this set of data points in the integral (1) gives a quantity $\text{Comp}(K)$ which is always finite in our main case of interest. Then the approximation of [7, 8, 15] is

$$\text{Comp}(K) = \frac{d}{2} \log \frac{n}{2\pi} + \log \int_K d\pi_\Theta(\theta) + o(1) \quad (2)$$

as $n \rightarrow \infty$, where d is the dimension of Θ , π_Θ is the Jeffreys prior on Θ (scaled so it does not depend on n) and $\int_K d\pi_\Theta(\theta) = \text{Vol}(K)$ is the Fisher information volume of K .

Suppose now that the data are IID and that \mathcal{M} is a natural exponential family, and let $\Theta \subseteq \mathbb{R}^d$ be any parameter space for \mathcal{M} , not necessarily the natural parameterisation. Assume that each distribution in \mathcal{M} satisfies Cramér's condition (which holds for all continuous distributions) and that $K \subseteq \Theta$ is a compact, d -dimensional subset of \mathbb{R}^d with smooth boundary (see the end of Section 2 for precise statements of these conditions). Then our main result is the following refinement of (2), for any non-negative integer s :

$$\text{Comp}(K) = \frac{d}{2} \log \frac{n}{2\pi} + \log \int_K \left(\sum_{i=0}^s F_i(\theta) n^{-i} \right) d\pi_\Theta(\theta) + o(n^{-s}) \quad (3)$$

as $n \rightarrow \infty$, where each $F_i(\theta)$ is a smooth function which does not depend on n .

Our methods can be used to calculate the asymptotic expansion (3) for $\text{Comp}(K)$ to any desired degree of accuracy. The first few of the functions $F_i(\theta)$ are given explicitly in Theorem 2 in terms of the Fisher information metric and cumulants. In particular, $F_0(\theta) = 1$, so taking $s = 0$ in (3) gives the approximation (2). Further, setting $s = 1$ in (3) shows, at least for the case of exponential families and IID data, that the approximation (2) is actually valid to order $O(n^{-1})$, which appears to be previously unknown.

The functions $F_i(\theta)$ which appear in (3) are likely to be interesting, since they arise naturally from the general, information-theoretic principles which underpin MDL. These functions appear to be more related to the statistical aspects of the model \mathcal{M} than to its curvature or other geometrical aspects, see Remark 4. Theorem 2 gives these functions in terms of the Fisher information metric and cumulants, but a more natural formulation in terms of the Amari-Chentsov tensors is also given in Section 7.

The rest of this paper is set out as follows. In Section 2 we define the main objects of this paper and fully describe the regularity conditions under which (3) holds. In Section

3 we give an asymptotic expansion for the parametric complexity in terms of a function which arises in Edgeworth expansions. Explicit formulae for the first few terms of this function are then obtained in Section 4, completing the proof of (3). In Section 5, the first two terms of (3) are calculated explicitly for some examples. In Section 6, the finite-sample performance of (2) and these higher order approximations are assessed, using an exact formula for the parametric complexity in the case of spherical normal data. We give alternative, co-ordinate-independent formulae for the functions $F_i(\theta)$ in Section 7, before finishing with a summary of the paper in Section 8. Appendix A contains the proofs of all of our results and Appendix B gives a fairly self-contained description of Edgeworth expansions and the Hermite numbers.

2 Definitions and regularity conditions

As in the Introduction, let \mathcal{M} be a natural exponential family of order d [13, §2.2.1] and let Θ be any parameter space of \mathcal{M} , which we assume without loss of generality is an open subset of \mathbb{R}^d . Then \mathcal{M} is a set of probability measures on \mathbb{R}^d , with each distribution in \mathcal{M} of the form $p_1(\cdot, \theta)\mu_1$ for some $\theta \in \Theta$, where: μ_1 is some measure on \mathbb{R}^d ; $p_1(\cdot, \theta)$ is the function $x \mapsto p_1(x, \theta)$ of x alone, with θ held fixed;

$$p_1(x, \theta) = \exp(x \cdot \eta(\theta) - \psi(\eta(\theta)));$$

$\eta : \Theta \rightarrow \mathbb{R}^d$ is the reparameterisation map from Θ to the natural parameter space (a diffeomorphism onto its image); ψ is the log-partition function, which is determined by the condition that each measure $p_1(\cdot, \theta)\mu_1$ is normalised; and $x \cdot \eta(\theta)$ is the Euclidean inner product of x and $\eta(\theta)$.

Let \mathcal{M}^n be the corresponding model for n IID observations $x_1, \dots, x_n \in \mathbb{R}^d$ and let $x^n = (x_1, \dots, x_n) \in \mathbb{R}^{nd}$ be the observed data for this model. Then \mathcal{M}^n is the set of all measures of the form $p_n(\cdot, \theta)\mu_n$ for $\theta \in \Theta$, where $p_n(x^n, \theta) = p_1(x_1, \theta) \dots p_1(x_n, \theta)$ is the product likelihood and $\mu_n = \mu_1 \times \dots \times \mu_1$ is the product measure on \mathbb{R}^{nd} .

The (scaled) Fisher information metric g_Θ on Θ is the $d \times d$ matrix-valued function on Θ with $(i, j)^{th}$ entry

$$g_{ij}(\theta) = -\frac{1}{n} \int \left[\frac{\partial^2 \log p_n(x^n, \theta)}{\partial \theta_i \partial \theta_j} \right] p_n(x^n, \theta) d\mu_n(x^n) \quad (4)$$

for $i, j = 1, \dots, d$, where the factor of $1/n$ ensures that g_Θ does not depend on n (since the data are IID). Then the Jeffreys prior π_Θ is $\sqrt{\det g_\Theta}$ times the Lebesgue measure on \mathbb{R}^d , and so is equal to the Fisher information volume density on Θ .

Let \mathcal{X}^n be the support of μ_n , which can be interpreted as the set of all possible values of the observed data x^n . Let $\hat{\theta} : \mathcal{X}^n \rightarrow \Theta$ be the maximum likelihood estimator for the parameter θ , so that $\hat{\theta}(x^n) \in \Theta$ is the maximum likelihood estimate corresponding to x^n . This estimate does not exist for all x^n but it is unique when it does exist, so $\hat{\theta}$ is well-defined if we restrict \mathcal{X}^n appropriately [2, Corollary 9.6]. If K is any compact subset of Θ then let

$$\hat{\theta}^{-1}(K) = \{x^n \in \mathcal{X}^n \mid \hat{\theta}(x^n) \in K\}$$

be the set of data points x^n whose maximum likelihood estimates exist and lie in K . Then define

$$\text{Comp}(K) = \log \int_{\hat{\theta}^{-1}(K)} p_n(x^n, \hat{\theta}(x^n)) d\mu_n(x^n) \quad (5)$$

to be the contribution to the parametric complexity corresponding to $K \subseteq \Theta$.

Now, the maximum likelihood estimate $t = \hat{\theta}(x^n) \in \Theta$ is a random variable, so write its distribution as $q_n(\cdot, \theta)\nu_n$, where $q_n(t, \theta)$ is some smooth function and ν_n is some measure on \mathbb{R}^d which is independent of θ . In fact, the family of distributions $q_n(\cdot, \theta)\nu_n$ for all $\theta \in \Theta$

is also an exponential family, though we will not need this result here. For now, we only use the fact that the maximum likelihood estimator is a sufficient statistic (by Theorem 2.2.6 of [13] and the paragraph preceding it) and that the parametric complexity can be calculated from the distribution of any sufficient statistic [5, §III.F]. Then (5) and the argument in [5, §III.F] imply

$$\text{Comp}(K) = \log \int_K q_n(t, t) d\nu_n(t), \quad (6)$$

which is intuitively reasonable since $q_n(t, t)$ is the maximized likelihood.

Remark 1. *Rissanen [15] proved (2) by using the central limit theorem to approximate the integrand of (6), despite the fact that the central limit theorem describes the distribution of t for any fixed θ while the integrand of (6) is $q_n(t, \theta(t))$, where $\theta(t) = t$ is not fixed. To get around this problem, Rissanen effectively replaced the function $\theta(t)$ by a step function $\bar{\theta}(t)$ which closely approximates $\theta(t)$, then he applied the central limit theorem in each of the small sets where $\bar{\theta}(t)$ is constant (and hence θ is fixed). By contrast, our approach will be to approximate the integral (6) by (a multiple of) the double integral of $q_n(t, \theta)$ with respect to t and θ , where (t, θ) ranges over a small neighbourhood of the set $\{(t, \theta) \mid \theta = t \text{ and } t \in K\}$. If we integrate with respect to t first and θ second in this double integral then θ is fixed in the innermost integral, so we can apply the central limit theorem (and higher-order Edgeworth expansions) there.*

We now impose the following regularity conditions, which will be needed when we use Edgeworth expansions in the next section to largely prove (3).

Condition 1. *We assume that each distribution Q in the exponential family \mathcal{M} satisfies Cramér's condition*

$$\limsup_{\|z\| \rightarrow \infty} |\hat{Q}(z)| < 1, \quad (7)$$

where \hat{Q} is the characteristic function of Q [6, eqn. 1.29]. Cramér's condition is satisfied for many interesting distributions, such as all continuous distributions (those which admit probability density functions) [12, p. 57] and the distributions of the minimal sufficient statistics of continuous exponential families [6, Lemma 1.5].

Condition 2. *In addition to assuming that $K \subseteq \Theta \subseteq \mathbb{R}^d$ is compact, we assume that K is a smooth d -manifold with boundary. This means that each point of K is at the centre of some small d -dimensional ball which intersects K either in the whole ball (for interior points of K) or in approximately a half ball (for boundary points of K).*

3 Higher order asymptotics

In this section, we will mostly work in the expectation parameter space of the exponential family \mathcal{M} , though the details of this parameter space are not important here beyond two specific properties, which we now recall. Let Ξ denote the expectation parameter space and let $\xi \in \Xi$ denote a generic expectation parameter (which is equal to the expected value of the sufficient statistic). If x_1, \dots, x_n are IID random variables governed by an unknown element of \mathcal{M} and $x^n = (x_1, \dots, x_n)$, as in Section 2, then the first property is that the maximum likelihood estimator $\hat{\xi}$ of the expectation parameter ξ is simply the mean, i.e.,

$$\hat{\xi}(x^n) = (x_1 + \dots + x_n)/n \quad (8)$$

by Theorem 2.2.6 of [13] and the comments preceding it. The second property is that this estimator is exactly (as opposed to asymptotically) unbiased and efficient, meaning that

$$\mathbb{E}[\hat{\xi}] = \xi \text{ and } \text{Var}(\hat{\xi}) = (ng_{\Xi}(\xi))^{-1} \quad (9)$$

by [13, Theorems 2.2.1 and 2.2.5], where g_{Ξ} is the Fisher information metric on Ξ and the superscripted -1 indicates the matrix inverse. The first property allows us to apply Edgeworth expansions directly to the maximum likelihood estimator $\hat{\xi}$ and the second property allows us to express the approximating distribution in terms of the Fisher information matrix.

Let $\Phi_{\xi} : \mathbb{R}^d \rightarrow \mathbb{R}$ be the probability density function (PDF) of the d -dimensional normal distribution $N_d(\xi, (ng_{\Xi}(\xi))^{-1})$ which has the same mean and variance as $\hat{\xi}$, so that

$$\Phi_{\xi}(x) = \left(\frac{n}{2\pi}\right)^{d/2} \sqrt{\det g_{\Xi}(\xi)} \exp\left(-\frac{n}{2}(x - \xi)^T g_{\Xi}(\xi)(x - \xi)\right) \quad (10)$$

for any $x \in \mathbb{R}^d$. As in Section 2 (but with $\xi, \hat{\xi}, \Xi, x$ in place of $\theta, \hat{\theta}, \Theta, t$), let the distribution of the maximum likelihood estimator $\hat{\xi}$ of ξ be $q_n(\cdot, \xi)\nu_n$. For any fixed $\xi \in \Xi$, (8) and the central limit theorem imply that this distribution can be approximated as

$$\int_B q_n(x, \xi) d\nu_n(x) = \int_B \Phi_{\xi}(x) d\lambda(x) + o(1) \quad (11)$$

as $n \rightarrow \infty$, where λ is the Lebesgue measure on \mathbb{R}^d and the approximation is uniform for all Borel sets $B \subseteq \Xi$ satisfying some weak regularity conditions (e.g. see [6, Corollary 1.4]). Edgeworth expansions [3, 6] refine this result to imply, for fixed ξ and for any integer or half-integer $s \geq 0$, that there is a function $H_s : \Xi \times \Xi \rightarrow \mathbb{R}$ so that

$$\int_B q_n(x, \xi) d\nu_n(x) = \int_B H_s(x, \xi) \Phi_{\xi}(x) d\lambda(x) + o(n^{-s}) \quad (12)$$

as $n \rightarrow \infty$, where the error is uniform for all Borel sets $B \subseteq \Xi$ satisfying weak regularity conditions, as for (11).

The function $H_s(x, \xi)$ appearing in (12) is a polynomial in x and only depends on ξ through the cumulants of the distribution in \mathcal{M} corresponding to ξ . In Section 4 we will give explicit formulae for the function $H_s(x, \xi)$ in the case of most interest to us. However, even without an explicit formula, the defining property (12) of $H_s(x, \xi)$ allows us to state and prove the following theorem, which is the main ingredient in our proof of (3).

Theorem 1. *Let \mathcal{M} and $K \subseteq \Theta$ satisfy the regularity conditions given at the end of Section 2, where Θ is any parameter space for \mathcal{M} . Let $f : \Theta \rightarrow \Xi$ be the reparameterisation map from Θ to the expectation parameter space Ξ . Then for any integer $s \geq 0$,*

$$\text{Comp}(K) = \frac{d}{2} \log \frac{n}{2\pi} + \log \int_K H_s(f(\theta), f(\theta)) d\pi_{\Theta}(\theta) + o(n^{-s}) \quad (13)$$

as $n \rightarrow \infty$.

Proof. See Section A.1 (and see Remark 1 for a heuristic description of the proof). \square

Remark 2. *Edgeworth expansions have a number of known weaknesses, such as potentially giving signed distributions (e.g., partly negative PDFs) and only controlling absolute errors, which are not very informative in the tails of distributions. However, neither of these defects affects our results, since we only need Edgeworth expansions in small neighbourhoods of the mean.*

4 The integrand on the natural parameter space

In this section, we complete the proof of (3) by calculating the integrand $H_s(f(\theta), f(\theta))$ appearing in (13) for the special case when Θ is the natural parameter space of the exponential family. Note that the corresponding function for any other parameter space

is simply obtained by composing the function found in this section with the appropriate reparameterisation map (see the first paragraph of Section A.1).

So let Θ be the natural parameter space and let $\psi : \Theta \rightarrow \mathbb{R}$ be the log-partition function of the natural exponential family \mathcal{M} , so that the distribution of one data point x is

$$e^{x \cdot \theta - \psi(\theta)} \mu_1 \quad (14)$$

for some measure μ_1 on \mathbb{R}^d which does not depend on θ . The cumulant generating function of the distribution (14) is $K(t) = \psi(t + \theta) - \psi(\theta)$ [13, eq. 2.2.4], so for any integer $r \geq 1$ and any $i_1, \dots, i_r \in \{1, \dots, d\}$, the cumulant $\kappa_{i_1 \dots i_r}$ of the distribution (14) is

$$\kappa_{i_1 \dots i_r} = \left. \frac{\partial^r K}{\partial t_{i_1} \dots \partial t_{i_r}} \right|_{t=0} = \left. \frac{\partial^r \psi}{\partial \theta_{i_1} \dots \partial \theta_{i_r}} \right|_{\theta=\theta}. \quad (15)$$

Let $g^{ij} = (g_\Theta(\theta)^{-1})_{ij}$ be the $(i, j)^{th}$ component of the matrix inverse of the Fisher information metric $g_\Theta(\theta)$ on Θ .

Theorem 2. *For Θ the natural parameter space, the integrand appearing in (13) is given by*

$$H_s(f(\theta), f(\theta)) = \sum_{i=0}^s F_i(\theta) n^{-i}$$

for any $\theta \in \Theta$, where each $F_i(\theta)$ is a function of θ but not n . In particular, $F_0(\theta) = 1$,

$$\begin{aligned} F_1(\theta) = & \frac{1}{8} \sum_{i_1, \dots, i_4} \kappa_{i_1 \dots i_4} g^{i_1 i_2} g^{i_3 i_4} - \frac{1}{8} \sum_{i_1, \dots, i_6} \kappa_{i_1 i_2 i_3} \kappa_{i_4 i_5 i_6} g^{i_1 i_2} g^{i_3 i_4} g^{i_5 i_6} \\ & - \frac{1}{12} \sum_{i_1, \dots, i_6} \kappa_{i_1 i_2 i_3} \kappa_{i_4 i_5 i_6} g^{i_1 i_4} g^{i_2 i_5} g^{i_3 i_6} \end{aligned} \quad (16)$$

and

$$\begin{aligned} F_2(\theta) = & \frac{1}{720} \sum_{i_1, \dots, i_6} \kappa_{i_1 \dots i_6} h_{i_1 \dots i_6} + \frac{1}{720} \sum_{i_1, \dots, i_8} \kappa_{i_1 i_2 i_3} \kappa_{i_4 \dots i_8} h_{i_1 \dots i_8} + \\ & \frac{1}{1152} \sum_{i_1, \dots, i_8} \kappa_{i_1 \dots i_4} \kappa_{i_5 \dots i_8} h_{i_1 \dots i_8} + \frac{1}{1728} \sum_{i_1, \dots, i_{10}} \kappa_{i_1 i_2 i_3} \kappa_{i_4 i_5 i_6} \kappa_{i_7 \dots i_{10}} h_{i_1 \dots i_{10}} + \\ & \frac{1}{31104} \sum_{i_1, \dots, i_{12}} \kappa_{i_1 i_2 i_3} \kappa_{i_4 i_5 i_6} \kappa_{i_7 i_8 i_9} \kappa_{i_{10} i_{11} i_{12}} h_{i_1 \dots i_{12}}, \end{aligned} \quad (17)$$

where $h_{i_1 \dots i_r}$ is defined below and each index i_k ranges from 1 to d in each sum, e.g.

$$\sum_{i_1, \dots, i_4} \text{ is shorthand for } \sum_{i_1=1}^d \sum_{i_2=1}^d \sum_{i_3=1}^d \sum_{i_4=1}^d.$$

Proof. The proof just rescales the Edgeworth expansions from Barndorff-Nielsen and Cox [3] and specialises them to exponential families, see Section A.2 for details. \square

The expressions $h_{i_1 \dots i_r}$ appearing in (17), where $r = 2k$ is an even positive integer, are the (un-normalized) Hermite numbers given by

$$h_{i_1 \dots i_r} = \frac{(-1)^k}{k!} \frac{\partial^r}{\partial z_{i_1} \dots \partial z_{i_r}} \left[\left(\sum_{a,b=1}^d g^{ab} z_a z_b / 2 \right)^k \right], \quad (18)$$

where z_1, \dots, z_d are dummy variables, see Theorem 5 in Appendix B. The expression in square brackets is a degree r polynomial in z_1, \dots, z_d so $h_{i_1 \dots i_r}$ does not depend on z_1, \dots, z_d . Up to sign, $h_{i_1 \dots i_r}$ is a certain sum of products of components g^{ab} of $g_\Theta(\theta)^{-1}$, e.g. see (39) and (40).

Remark 3. *The un-normalized Hermite numbers (which corresponding to a normal distribution with a general, rather than identity, variance-covariance matrix) also arise in Gaussian processes and quantum field theory [14, eqn. 1 and 2].*

Remark 4. *The geometrical or statistical significance of the functions $F_1(\theta), F_2(\theta), \dots$ is not clear. $F_1(\theta)$ does not appear to be any sort of contraction of the Riemann curvature tensor R (for either the Levi-Civita connection or any of Amari's other α -connections [1, §2.3]) because in the natural parameter space, R only involves third derivatives of the log-partition function ψ whereas $F_1(\theta)$ involves fourth derivatives. Some of the terms in (16) have a superficial similarity to Efron's curvature (see [9] or [13, eqn. 4.3.10]), but Efron's formula measures the extrinsic curvature of curved exponential families and so is 0 for exponential families. In Section 7, we give an expression for $F_1(\theta)$ that is valid in any parameter space, in terms of contractions of products of the Amari-Chentsov tensors.*

5 Examples

It is possible to calculate the functions $F_i(\theta)$ appearing in (3) by using the formulae of Theorem 2 and a formula-manipulation program like Maxima [18]. This is especially easy when d is small, though it can also be done for all d , as we now illustrate.

5.1 Exponential data

Let \mathcal{M} be the family of exponential distributions on \mathbb{R} , so each observation x is governed by some distribution of the form (14), where μ_1 is the Lebesgue measure restricted to the positive reals, the natural parameter space Θ is the set of negative reals, $-\theta > 0$ is the rate parameter of the exponential distribution and the log-partition function is

$$\psi(\theta) = -\log(-\theta).$$

Each distribution in \mathcal{M} is continuous so it satisfies Cramér's condition by [12, p. 57].

Since the parameter space Θ of this family has dimension $d = 1$, the indices i_k in (16) and (17) all take the value $i_k = 1$, so each sum consists of a single term. Also, g_{11} is the second derivative of $\psi(\theta)$ by (4), so $g_{11} = \theta^{-2}$, and the matrix inverse of the 1×1 matrix g_Θ has the single component $g^{11} = \theta^2$. Also, by (15), the cumulants appearing in Theorem 2 are just the derivatives of $\psi(\theta)$, e.g. $\kappa_{111} = -2\theta^{-3}$ and $\kappa_{1111} = 6\theta^{-4}$. So by (16),

$$\begin{aligned} F_1(\theta) &= \frac{1}{8}\kappa_{1111}g^{11}g^{11} - \frac{1}{8}\kappa_{111}\kappa_{111}g^{11}g^{11}g^{11} - \frac{1}{12}\kappa_{111}\kappa_{111}g^{11}g^{11}g^{11} \\ &= \frac{1}{8}(6\theta^{-4})(\theta^2)^2 - \frac{1}{8}(-2\theta^{-3})^2(\theta^2)^3 - \frac{1}{12}(-2\theta^{-3})^2(\theta^2)^3 \\ &= -1/12. \end{aligned}$$

A similar calculation, based on (17) and aided by Maxima [18], shows that

$$F_2(\theta) = 1/288.$$

Since we have just shown that $F_1(\theta)$ and $F_2(\theta)$ are constant in θ , we will now write F_1 and F_2 for these constant functions. Then (3) implies

$$\begin{aligned} \text{Comp}(K) &= \frac{d}{2} \log \frac{n}{2\pi} + \log \int_K (1 + F_1 n^{-1} + F_2 n^{-2}) d\pi_\Theta(\theta) + O(n^{-3}) \\ &= \frac{d}{2} \log \frac{n}{2\pi} + \log \int_K d\pi_\Theta(\theta) + \log (1 + F_1 n^{-1} + F_2 n^{-2}) + O(n^{-3}) \\ &= \frac{d}{2} \log \frac{n}{2\pi} + \log \text{Vol}(K) + (F_1 n^{-1} + F_2 n^{-2}) - \frac{1}{2} (F_1 n^{-1} + F_2 n^{-2})^2 + O(n^{-3}) \\ &= \frac{d}{2} \log \frac{n}{2\pi} + \log \text{Vol}(K) + F_1 n^{-1} + (F_2 - F_1^2/2) n^{-2} + O(n^{-3}), \end{aligned} \tag{19}$$

where the second last step uses the asymptotic expansion $\log(1+z) = z - z^2/2 + O(z^3)$ as $z \rightarrow 0$. Substituting the above values of F_1 and F_2 into (19) therefore gives

$$\text{Comp}(K) = \frac{d}{2} \log \frac{n}{2\pi} + \log \text{Vol}(K) - \frac{1}{12n} + O(n^{-3})$$

because $F_2 - F_1^2/2$ vanishes.

Note that (19) holds for general d and even for non-constant $F_1(\theta)$ and $F_2(\theta)$ if we replace each F_i in (19) by the expected value of $F_i(\theta)$, where θ is regarded as a random variable whose distribution is proportional to the Jeffreys prior restricted to K .

5.2 Spherical normal data with unknown variance

We now consider a model where each observation y is distributed according to the $(d-1)$ -dimensional spherical normal distribution $N_{d-1}(\beta, \sigma^2 I_{d-1})$, where I_{d-1} is the $(d-1) \times (d-1)$ identity matrix and the parameters $\beta \in \mathbb{R}^{d-1}$ and $\sigma > 0$ are to be estimated. One reason for studying this model is that an exact formula for $\text{Comp}(K)$ is known, so in Section 6 we will verify the expression for $F_1(\theta)$ given here.

For this model, a sufficient statistic x and the corresponding natural parameter θ are given by

$$x = \begin{bmatrix} y \\ \|y\|^2 \end{bmatrix} \text{ and } \theta = \frac{1}{\sigma^2} \begin{bmatrix} \beta \\ -1/2 \end{bmatrix}, \quad (20)$$

where $y = [y_1 \dots y_{d-1}]^T$ and $\|y\|^2 = y_1^2 + \dots + y_{d-1}^2$. Then x is governed by a member of a natural exponential family of the form (14), with log-partition function

$$\psi(\theta) = \left(\frac{1-d}{2} \right) \log(-2\theta_d) - (\theta_1^2 + \dots + \theta_{d-1}^2)/4\theta_d.$$

Note that Cramér's condition holds for each distribution in this natural exponential family, by [6, Lemma 1.5]. The (scaled) Fisher information metric $g_\Theta(\theta)$ is the Hessian of $\psi(\theta)$, by (4) or [13, Theorem 2.2.5], and the matrix inverse of g_Θ has $(i, j)^{th}$ component

$$g^{ij} = \begin{cases} \theta_i \theta_j (d-1)/2 - 2\theta_d & \text{if } i = j \neq d \\ \theta_i \theta_j (d-1)/2 & \text{otherwise.} \end{cases} \quad (21)$$

Also, the cumulants $\kappa_{i_1 \dots i_r}$ can be obtained by differentiating ψ , as in (15). Then a calculation based on (16) (which was aided by Maxima [18] and which considered all combinations of the two cases $i_k < d$ and $i_k = d$ for each index i_k) shows that the three sums appearing in (16) are constant in θ and are equal to the expressions in square brackets below:

$$\begin{aligned} F_1(\theta) &= \frac{1}{8} \left[\frac{8d+4}{d-1} \right] - \frac{1}{8} \left[\frac{2d^2+4d+2}{d-1} \right] - \frac{1}{12} \left[\frac{6d+2}{d-1} \right] \\ &= \frac{1-3d^2}{12(d-1)}. \end{aligned} \quad (22)$$

So when $s = 1$, (3) becomes

$$\begin{aligned} \text{Comp}(K) &= \frac{d}{2} \log \frac{n}{2\pi} + \log \int_K \left[1 + \frac{1-3d^2}{12(d-1)n} \right] d\pi_\Theta(\theta) + o(n^{-1}) \\ &= \frac{d}{2} \log \frac{n}{2\pi} + \log \int_K d\pi_\Theta(\theta) + \log \left[1 + \frac{1-3d^2}{12(d-1)n} \right] + o(n^{-1}) \\ &= \frac{d}{2} \log \frac{n}{2\pi} + \log \text{Vol}(K) + \frac{1-3d^2}{12(d-1)n} + o(n^{-1}) \end{aligned} \quad (23)$$

as $n \rightarrow \infty$. Note that this is just the usual approximation (2) plus the correction term $(1-3d^2)/12(d-1)n$. We will investigate the finite-sample performance of the two approximations (23) and (2) in Section 6, below.

5.3 Spherical normal data with known variance

For spherical normal data with known variance, the log-partition function $\psi(\theta)$ is quadratic in θ , so all of the cumulants $\kappa_{i_1 \dots i_r}$ vanish for $r \geq 3$, and hence the functions $F_i(\theta)$ also vanish for $i \geq 1$. So by (3), the standard formula (2) is accurate to order $o(n^{-s})$ for all integers $s \geq 0$. This is a positive check on our formulae, because it is well-known that (2) is the exact formula for this model, as we now briefly demonstrate.

Working in the expectation parameter space, if $x_1, \dots, x_n \sim N_d(\xi, \sigma^2 I_d)$ for $\sigma > 0$ known and $\xi \in \mathbb{R}^d$ unknown, then $\hat{\xi} = (x_1 + \dots + x_n)/n$ and hence $\hat{\xi} \sim N_d(\xi, (\sigma^2/n)I_d)$. Therefore $\nu_n = \lambda$ and

$$q_n(x, \xi) = \left(\frac{n}{2\pi\sigma^2}\right)^{d/2} \exp\left(-\frac{n}{2\sigma^2}\|x - \xi\|^2\right)$$

so $q_n(\xi, \xi) = (n/2\pi\sigma^2)^{d/2}$. Hence if $K \subseteq \Xi$ then the exact formula for $\text{Comp}(K)$ is

$$\text{Comp}(K) = \log \int_K q_n(\xi, \xi) d\nu_n(\xi) = \log \int_K (n/2\pi\sigma^2)^{d/2} d\lambda(\xi) = \frac{d}{2} \log \frac{n}{2\pi} + \log \text{Vol}(K)$$

since $\pi_\Xi = \lambda\sigma^{-d}$.

6 Finite sample performance

In this section we compare the finite-sample performance of the standard approximation (2) and its first-order correction (i.e., (3) with $s = 1$), in the case of the spherical normal model with unknown variance (see Section 5.2).

Rissanen [16] studied the linear regression model with the parameterisation $\theta = (\beta, \tau)$, where $\tau = \sigma^2$ is the noise variance and β is the regression coefficient (note that this is not the natural parameterisation of this exponential family). Let $t = (\hat{\beta}, \hat{\tau})$ be the maximum likelihood estimate for θ . If the model has N observations, k covariates and design matrix A (i.e., A is the $N \times k$ matrix whose columns are the covariates of the model) then Rissanen showed that the maximized likelihood is

$$q(t, t) = \frac{\sqrt{\det A^T A}}{(\pi N)^{k/2} \Gamma(\frac{N-k}{2})} \left(\frac{N}{2e}\right)^{N/2} \hat{\tau}^{-1-k/2} \quad (24)$$

with respect to the Lebesgue measure λ , where Γ is the gamma function. It is not hard to show that the unscaled Jeffrey's prior $\bar{\pi}_\Theta$ for the above parameterisation is given by

$$d\bar{\pi}_\Theta(\theta) = \tau^{-1-k/2} \sqrt{(N/2) \det A^T A} d\lambda(\theta), \quad (25)$$

where ‘unscaled’ means that $\bar{\pi}_\Theta$ is the volume density of the unscaled Fisher information metric, which is n times the right-hand side of (4). Therefore

$$\begin{aligned} \text{Comp}(K) &= \log \int_K q(t, t) d\lambda(t) \text{ by (6)} \\ &= \log \int_K q(\theta, \theta) d\lambda(\theta) \text{ since } t \text{ is a dummy variable} \\ &= \log \int_K \frac{(\pi N)^{-k/2}}{\Gamma(\frac{N-k}{2})} \sqrt{\frac{2}{N}} \left(\frac{N}{2e}\right)^{N/2} d\bar{\pi}_\Theta(\theta) \text{ by (24) and (25)} \\ &= \log \int_K d\bar{\pi}_\Theta(\theta) + \log \left(\frac{(\pi N)^{-k/2}}{\Gamma(\frac{N-k}{2})} \sqrt{\frac{2}{N}} \left(\frac{N}{2e}\right)^{N/2} \right) \end{aligned} \quad (26)$$

where the last step follows because the integrand is constant in θ . Note that this is not an approximation, $\text{Comp}(K)$ is given exactly by (26).

Now, take A to be the $N \times k$ matrix $A = [I_k \cdots I_k]^T$, where $k = d - 1$, $N = nk$ and I_k is the $k \times k$ identity matrix. Then the regression response variable $[y_1 \cdots y_n]^T$ is distributed in such a way that $y_1, \dots, y_n \sim N_{d-1}(\beta, \sigma^2 I_{d-1})$ are IID, as in Section 5.2. Substituting these values into (26) therefore gives the following exact formula for the parametric complexity $\text{Comp}(K)$ of the model of Section 5.2:

$$\text{Comp}(K) = \log \text{Vol}(K) - \frac{k}{2} \log(\pi k) - \frac{1}{2} \log \frac{k}{2} + \frac{nk}{2} \log \frac{nk}{2e} - \log \Gamma\left(\frac{k(n-1)}{2}\right), \quad (27)$$

where we have retained the notation $k = d - 1$ for convenience and where, as is standard throughout this paper, $\text{Vol}(K) = \int_K d\pi_\Theta(\theta) = n^{-d/2} \int_K d\bar{\pi}_\Theta(\theta)$ is the volume of K with respect to the scaled Jeffreys prior π_Θ . We can use this exact formula in two ways: to assess the accuracy of our asymptotic expansion and to check the veracity of our calculations.

Figure 1 assesses the accuracy of both the standard approximation (2) to $\text{Comp}(K)$, in the case of spherical normal data with unknown variance, and its first-order correction (23). Note that the summand $\log \text{Vol}(K)$ cancels when comparing these two approximations to the exact formula (27), so the amount of over- or underestimation of $\text{Comp}(K)$ does not depend on K . Figure 1 shows that both formulae overestimate $\text{Comp}(K)$, and hence NML codelength, for all model dimensions d and all sample sizes $n \geq d$, with the amount of overestimation increasing with d and decreasing with n (note that results cannot be calculated for the under-determined models with $n < d$.) The overestimation is negligible for the corrected formula (23) when $n \geq 20$, but it is significant for the standard formula (2) for moderate n , e.g. it is greater than 0.05 for $n = 50$ and $d = 11$.

Substituting an asymptotic expansion for the gamma function into (27) gives

$$\text{Comp}(K) = \frac{d}{2} \log \frac{n}{2\pi} + \log \text{Vol}(K) - n^{-1} \left[\frac{3d^2 - 1}{12(d-1)} \right] - n^{-2} \left[\frac{d(d+1)}{12(d-1)} \right] + O(n^{-3}).$$

Comparing the above formula to (23), which was obtained directly from Theorems 1 and 2, shows that our expression (22) for $F_1(\theta)$ is correct. The coefficient of n^{-2} in the above formula also verifies our calculation (not shown here) of $F_2(\theta)$ for $d = 2$. This is therefore a positive check on the formulae of Theorems 1 and 2.

7 The integrand on a general parameter space

In this section, we give an expression for the integrand $H_s(f(\theta), f(\theta))$ of (13) which is valid in all parameter spaces. The expression for $H_s(f(\theta), f(\theta))$ in Theorem 2 is convenient for calculations, but it is not invariant under co-ordinate changes because cumulants are not tensors on Θ . The cumulant for a distribution is independent of the family to which it belongs (cumulants are derivatives of $K(t)$ with respect to the dummy variable t) so the cumulants $\kappa_{i_1 \dots i_r}$ transform like functions rather than the components of a tensor. A co-ordinate-dependent expression like the one in Theorem 2 is likely to be arbitrary in some ways, and therefore likely to hide the true mathematical structure, so in this section we give an expression for $F_1(\theta)$ which holds for all parameterisations.

For any $r = 1, 2, 3, \dots$, define the r^{th} Amari-Chentsov tensor $\text{AC}^{(r)}$ to be the symmetric tensor on Θ with components

$$\text{AC}_{i_1, \dots, i_r}^{(r)} = \int \left(\prod_{j=1}^r \frac{\partial \log p_1(x, \theta)}{\partial \theta_{i_j}} \right) p_1(x, \theta) d\mu_1(x) \quad (28)$$

where $p_1(x, \theta)$ is the likelihood function at parameter θ for one observation x . This is a tensor on any parameter space because it has a co-ordinate-independent definition (essentially just by replacing the partial differentials with arbitrary first-order partial differential operators, which are interpreted as vector fields in a way which is standard in differential geometry).




Figure 1: The amount of overestimation of $\text{Comp}(K)$, and hence NML codelength, by both the standard approximation (2), corresponding to $s = 0$ in (3), and its first-order correction, corresponding to $s = 1$ in (3), as a function of sample size n and model dimension d , for the spherical normal model with unknown variance (see Section 6).

It is easy to show that the first Amari-Chentsov tensor $\text{AC}^{(1)}$ is trivial, $\text{AC}^{(1)} = 0$, but the second one is the Fisher information metric, $\text{AC}^{(2)} = g_{\Theta}$, so the Amari-Chentsov tensors can be seen as generalisations of the Fisher information metric tensor. The third Amari-Chentsov tensor $\text{AC}^{(3)}$ is also interesting, because the difference between any two of Amari's α -connections is a multiple of $\text{AC}^{(3)}$ [1, §2.3].

In the natural parameterisation of an exponential family, the Amari-Chentsov tensors can be calculated using the following recurrence relation.

Lemma 3. *If Θ is the natural parameterisation of the exponential family and $r \geq 3$ is an integer then*

$$\text{AC}_{i_1, \dots, i_r}^{(r)} = \frac{\partial}{\partial \theta_{i_r}} \text{AC}_{i_1, \dots, i_{r-1}}^{(r-1)} + \sum_{j=1}^{r-1} g_{i_j i_r} \text{AC}_{i_1, \dots, \hat{i}_j, \dots, i_{r-1}}^{(r-2)}$$

where the hat indicates that the term i_j should be left out.

Proof. See Section A.3. □

So in the natural parameterisation, using $\text{AC}_i^{(1)} = 0$, $\text{AC}_{ij}^{(2)} = g_{ij} = \partial^2 \psi / \partial \theta_i \partial \theta_j$ and Lemma 3, we have

$$\text{AC}_{i_1 i_2 i_3}^{(3)} = \frac{\partial^3 \psi}{\partial \theta_{i_1} \partial \theta_{i_2} \partial \theta_{i_3}}, \quad (29)$$

so another application of Lemma 3 implies

$$\text{AC}_{i_1 \dots i_4}^{(4)} = \frac{\partial^4 \psi}{\partial \theta_{i_1} \dots \partial \theta_{i_4}} + g_{i_1 i_2} g_{i_3 i_4} + g_{i_1 i_3} g_{i_2 i_4} + g_{i_1 i_4} g_{i_2 i_3}. \quad (30)$$

The above derivatives of ψ can be identified with the cumulants, by (15), so combining (29) and (30) with the expression for $F_1(\theta)$ in (16) gives

$$F_1(\theta) = \frac{1}{8} \sum_{i_1, \dots, i_4} \text{AC}_{i_1 \dots i_4}^{(4)} g^{i_1 i_2} g^{i_3 i_4} - \frac{1}{8} \sum_{i_1, \dots, i_6} \text{AC}_{i_1 i_2 i_3}^{(3)} \text{AC}_{i_4 i_5 i_6}^{(3)} g^{i_1 i_2} g^{i_3 i_4} g^{i_5 i_6} \\ - \frac{1}{12} \sum_{i_1, \dots, i_6} \text{AC}_{i_1 i_2 i_3}^{(3)} \text{AC}_{i_4 i_5 i_6}^{(3)} g^{i_1 i_4} g^{i_2 i_5} g^{i_3 i_6} - \frac{d^2 + 2d}{8} \quad (31)$$

in the natural parameterisation, where the last term arises from contractions between the metric and its inverse, for example

$$\sum_{i_1, \dots, i_4} g_{i_1 i_2} g_{i_3 i_4} g^{i_1 i_2} g^{i_3 i_4} = d^2.$$

But since the left- and right-hand sides of (31) are equivariant (i.e., transform correctly) under reparameterisations, and since we have just shown that (31) is valid for the natural parameterisation, (31) must hold for all parameterisations.

Using similar methods, an expression can also be given for $F_2(\theta)$, and higher terms, which is invariant under co-ordinate changes on the parameter space.

8 Summary

We have given an asymptotic expansion for the contribution $\text{Comp}(K)$ to the parametric complexity coming from a given set K in the parameter space, in the case of exponential families and IID data. Our methods allow this expansion to be calculated to any desired degree of accuracy, and we gave explicit formulae for the first three terms. The first term in this expansion is Rissanen's $o(1)$ approximation (2), so our proof specialises to the first demonstration that Rissanen's result is actually $O(n^{-1})$. Our proof is also arguably simpler than Rissanen's, since we do not have to partition K into optimally sized pieces and since the theory of Edgeworth expansions largely takes care of effects at the boundary of K for us. We calculated the first-order correction to Rissanen's approximation for some examples, and showed that the resulting expression has improved finite-sample behaviour. We also discussed the meaning of the higher-order terms, and we gave an expression for them which is valid in all parameterisations, in addition to giving a simpler expression in terms of cumulants which is valid for the natural parameterisation.

A Proofs

A.1 Proof of Theorem 1

Under the reparameterisation map $f : \Theta \rightarrow \Xi$, the measure π_Θ corresponds to π_Ξ and the function $\theta \mapsto H_s(f(\theta), f(\theta))$ corresponds to $\xi \mapsto H_s(\xi, \xi)$, so

$$\int_K H_s(f(\theta), f(\theta)) d\pi_\Theta(\theta) = \int_{f(K)} H_s(\xi, \xi) d\pi_\Xi(\xi).$$

Also, if we interpret K as a subset of Θ and $f(K)$ as a subset of Ξ , then

$$\text{Comp}(K) = \text{Comp}(f(K))$$

by the definition (5), since $\hat{\xi} = f \circ \hat{\theta}$ so the set of data points x^n with $\hat{\theta}(x^n) \in K$ is the same as the set of x^n with $\hat{\xi}(x^n) \in f(K)$. Therefore the general case of the theorem will follow from the special case where $\Theta = \Xi$ and f is the identity map, so we restrict to this case now.

The Riemannian metric g_Ξ can be used in a standard way to define a distance $d_F(x, \xi)$ between any two points $x, \xi \in \Xi$, so that $d_F(x, \xi)$ is equal to the infimum of the lengths of all smooth paths joining x and ξ . Then given any $x \in \Xi$ and any $\delta > 0$, let

$$B(x, \delta) = \{\xi \in \Xi \mid d_F(x, \xi) \leq \delta\}$$

be the closed ball in Ξ of radius δ centred at x . We assume that δ is small enough that $B(x, \delta)$ is compact for every $x \in K$ (i.e., that δ is less than the distance from K to the ‘boundary’ of Ξ).

Let v_δ be the volume of a Euclidean ball of radius δ . We will prove the theorem in the following three steps:

$$\begin{aligned} \exp(\text{Comp}(K)) &= \int_K \int_{B(x, \delta)} v_\delta^{-1} q_n(x, \xi) d\pi_\Xi(\xi) d\nu_n(x) + O(\delta) \text{ Step 1} \\ &= \int_K \int_{B(x, \delta)} v_\delta^{-1} H_s(x, \xi) \Phi_\xi(x) d\pi_\Xi(\xi) d\lambda(x) + O(\delta) + o(n^{-s}) \text{ Step 2} \\ &= \left(\frac{n}{2\pi}\right)^{d/2} \int_K H_s(\xi, \xi) d\pi_\Xi(\xi) + O(\delta) + o(n^{-s}) \text{ Step 3} \end{aligned}$$

as $n \rightarrow \infty$ and $\delta \rightarrow 0$. So taking the logarithm of both sides and setting $\delta = o(n^{-s})$ proves the theorem, since we will show that the error term $o(n^{-s})$ arising from Step 2 is uniform in δ .

Steps 1 and 3 are similar, and are proved below by approximating the integrand to zeroth order in δ then integrating out the inner integral in the double integral (i.e., integrating in the ξ direction). Step 2 is obtained by swapping the order of integration then using Edgeworth expansions in the inner integral (since ξ is fixed there).

Step 1. Let

$$R = \{(x, \xi) \in \Xi \times \Xi \mid x \in K \text{ and } \xi \in B(x, \delta)\}$$

be the region of integration in the double integrals above. Then for any (x, ξ) in R , the zeroth order Taylor series expansion of $\xi \mapsto q_n(x, \xi)$ about x shows that

$$q_n(x, \xi) = q_n(x, x) + O(\delta), \quad (32)$$

where the error term $O(\delta)$ is uniformly bounded for all $(x, \xi) \in R$ because $q_n(x, \xi)$ is a smooth function and R is compact. Also, the Fisher information volume of the ball $B(x, \delta)$ can be approximated by the volume v_δ of a Euclidean ball of the same radius

$$\int_{B(x, \delta)} d\pi_\Xi(\xi) = v_\delta + O(\delta) \quad (33)$$

by Gray [10], where the error term is uniformly bounded for all $x \in K$. Therefore

$$\begin{aligned} &\int_K \int_{B(x, \delta)} v_\delta^{-1} q_n(x, \xi) d\pi_\Xi(\xi) d\nu_n(x) \\ &= \int_K \int_{B(x, \delta)} v_\delta^{-1} q_n(x, x) d\pi_\Xi(\xi) d\nu_n(x) + O(\delta) \text{ by (32)} \\ &= \int_K v_\delta^{-1} q_n(x, x) \left(\int_{B(x, \delta)} d\pi_\Xi(\xi) \right) d\nu_n(x) + O(\delta) \\ &= \int_K q_n(x, x) d\nu_n(x) + O(\delta) \text{ by (33)} \\ &= \exp(\text{Comp}(K)) + O(\delta) \text{ by (6)} \end{aligned}$$

which proves Step 1.

Step 2. Define $C(\xi, \delta)$ to be the truncated ball $K \cap B(\xi, \delta)$ and define $N(K, \delta)$ to be the closed δ -neighbourhood of K in Ξ . Then

$$x \in K \text{ and } \xi \in B(x, \delta) \Leftrightarrow \xi \in N(K, \delta) \text{ and } x \in C(\xi, \delta) \quad (34)$$

so the region of integration R consists of all pairs $(x, \xi) \in \Xi \times \Xi$ satisfying either of these conditions. Therefore

$$\begin{aligned} & \int_K \int_{B(x, \delta)} v_\delta^{-1} q_n(x, \xi) d\pi_\Xi(\xi) d\nu_n(x) \\ &= \int_{N(K, \delta)} \int_{C(\xi, \delta)} v_\delta^{-1} q_n(x, \xi) d\nu_n(x) d\pi_\Xi(\xi) \text{ by (34) and Fubini's theorem} \\ &= \int_{N(K, \delta)} \int_{C(\xi, \delta)} v_\delta^{-1} H_s(x, \xi) \Phi_\xi(x) d\lambda(x) d\pi_\Xi(\xi) + o(n^{-s}) \text{ by (12)} \\ &= \int_K \int_{B(x, \delta)} v_\delta^{-1} H_s(x, \xi) \Phi_\xi(x) d\pi_\Xi(\xi) d\lambda(x) + o(n^{-s}) \text{ by (34) and Fubini's theorem} \end{aligned}$$

so Step 2 is proved. The fact that the error term $o(n^{-s})$ above is uniform in δ follows from [6, Corollary 1.4] with $\mathcal{A} = \{K \cap B(\xi, \delta) \mid \xi \in \Xi \text{ and } \delta > 0\}$, where K is fixed.

Step 3. For any (x, ξ) in the region of integration R , the zeroth order Taylor series expansion of $\xi \mapsto H_s(x, \xi) \Phi_\xi(x)$ about x shows that

$$\begin{aligned} H_s(x, \xi) \Phi_\xi(x) &= H_s(x, x) \Phi_x(x) + O(\delta) \\ &= \left(\frac{n}{2\pi}\right)^{d/2} H_s(x, x) \sqrt{\det g_\Xi(x)} + O(\delta), \end{aligned} \quad (35)$$

where the second equality follows by the definition (10) of Φ_ξ . Here the error term $O(\delta)$ is uniformly bounded for $(x, \xi) \in R$, since $H_s(x, \xi) \Phi_\xi(x)$ is a smooth function and R is compact. Therefore

$$\begin{aligned} & \int_K \int_{B(x, \delta)} v_\delta^{-1} H_s(x, \xi) \Phi_\xi(x) d\pi_\Xi(\xi) d\lambda(x) \\ &= \left(\frac{n}{2\pi}\right)^{d/2} \int_K \int_{B(x, \delta)} v_\delta^{-1} H_s(x, x) \sqrt{\det g_\Xi(x)} d\pi_\Xi(\xi) d\lambda(x) + O(\delta) \text{ by (35)} \\ &= \left(\frac{n}{2\pi}\right)^{d/2} \int_K v_\delta^{-1} H_s(x, x) \sqrt{\det g_\Xi(x)} \left(\int_{B(x, \delta)} d\pi_\Xi(\xi) \right) d\lambda(x) + O(\delta) \\ &= \left(\frac{n}{2\pi}\right)^{d/2} \int_K H_s(x, x) \sqrt{\det g_\Xi(x)} d\lambda(x) + O(\delta) \text{ by (33)} \\ &= \left(\frac{n}{2\pi}\right)^{d/2} \int_K H_s(x, x) d\pi_\Xi(x) + O(\delta) \\ &= \left(\frac{n}{2\pi}\right)^{d/2} \int_K H_s(\xi, \xi) d\pi_\Xi(\xi) + O(\delta) \end{aligned}$$

since $\pi_\Xi = \lambda \sqrt{\det g_\Xi}$ by definition of the Jeffreys prior and since x and ξ are dummy variables in the last two integrals. This proves Step 3 and hence the theorem.

A.2 Proof of Theorem 2

Recall that x_1, \dots, x_n are IID random variables governed by an unknown element of the exponential family \mathcal{M} . Suppose that the data-generating distribution has expectation parameter $\xi \in \Xi$, and consider this to be fixed throughout this section. The maximum likelihood estimate $\hat{\xi}$ of ξ is simply the sample mean $(x_1 + \dots + x_n)/n$, by (8), and its mean and variance are $\mathbb{E}[\hat{\xi}] = \xi$ and $\text{Var}(\hat{\xi}) = (ng_\Xi(\xi))^{-1}$, by (9). Then the Edgeworth

expansions from Barndorff-Nielsen and Cox [3] or Kass and Vos [13, §4.5] approximate the distribution ζ_n of the transformed mean $Z_n = \sqrt{n}(\hat{\xi} - \xi)$ by the asymptotic expansion

$$\int_B d\zeta_n(z) = \int_B \tilde{H}_s(z) \phi(z) d\lambda(z) + o(n^{-s}) \quad (36)$$

as $n \rightarrow \infty$, where $\tilde{H}_s(z)$ is an explicit polynomial in z , ϕ is the PDF for $N_d(0, g_\Xi(\xi)^{-1})$ and B is any Borel set satisfying some mild regularity conditions.

To find $H_s(x, \xi)$ in terms of $\tilde{H}_s(z)$, we simply have to calculate the effect of (the inverse of) the transformation $\tau : x \mapsto \sqrt{n}(x - \xi)$ on (36). The transformation τ takes $\hat{\xi}$ to Z_n , so τ^{-1} takes the distribution in the left-hand side of (36) to the distribution $dq_n(x, \xi) d\nu_n(x)$ of $\hat{\xi}$. On the other hand, the change of variables formula for PDFs shows that if W is any continuous random variable and ρ is the PDF of $\tau(W)$ then the PDF for W is $J_\tau \rho \circ \tau$, where $J_\tau = n^{d/2}$ is the Jacobian determinant of τ . Therefore τ^{-1} takes the distribution in the right-hand side of (36) to $\tilde{H}_s(\sqrt{n}(x - \xi)) \Phi_\xi(x) d\lambda(x)$, so (36) transforms under τ^{-1} to

$$\int_B dq_n(x, \xi) d\nu_n(x) = \int_B \tilde{H}_s(\sqrt{n}(x - \xi)) \Phi_\xi(x) d\lambda(x) + o(n^{-s}).$$

Comparing this equation to (12) gives $H_s(x, \xi) = \tilde{H}_s(\sqrt{n}(x - \xi))$, so

$$H_s(\xi, \xi) = \tilde{H}_s(0). \quad (37)$$

Now, the function $\tilde{H}_s(z)$ of [3] or [13, §4.5] is a sum of terms involving integral and half-integral powers of n , but all of the terms involving half-integral powers vanish at $z = 0$ by the theorem in the Appendix of [3]. So (37) implies that

$$H_s(f(\theta), f(\theta)) = \sum_{i=0}^s F_i(\theta) n^{-i}$$

for integral s , as claimed. The explicit formulae of [13, §4.5] then show that $F_0(\theta) = 1$ and

$$F_1(\theta) = \frac{1}{24} \sum_{i_1, \dots, i_4} \kappa_{i_1 \dots i_4} h_{i_1 \dots i_4} + \frac{1}{72} \sum_{i_1, \dots, i_6} \kappa_{i_1 i_2 i_3} \kappa_{i_4 i_5 i_6} h_{i_1 \dots i_6}. \quad (38)$$

The methods of [13, §4.5] allow any $F_i(\theta)$ to be calculated, for example see Appendix B, where the formula (17) for $F_2(\theta)$ is calculated explicitly. Therefore the theorem will be proved if we can convert (38) to the formula for $F_1(\theta)$ which appears in the statement of the theorem.

From [13, §4.5] or directly from the definition (18), the first few Hermite numbers are $h_{i_1 i_2} = -g^{i_1 i_2}$,

$$h_{i_1 \dots i_4} = g^{i_1 i_2} g^{i_3 i_4} + g^{i_1 i_3} g^{i_2 i_4} + g^{i_1 i_4} g^{i_2 i_3} \quad (39)$$

and

$$\begin{aligned} h_{i_1 \dots i_6} = & -(g^{i_1 i_2} g^{i_3 i_4} g^{i_5 i_6} + g^{i_1 i_3} g^{i_2 i_4} g^{i_5 i_6} + g^{i_1 i_4} g^{i_2 i_3} g^{i_5 i_6} + g^{i_1 i_2} g^{i_3 i_5} g^{i_4 i_6} + \\ & g^{i_1 i_3} g^{i_2 i_5} g^{i_4 i_6} + g^{i_1 i_5} g^{i_2 i_3} g^{i_4 i_6} + g^{i_1 i_2} g^{i_3 i_6} g^{i_4 i_5} + g^{i_1 i_3} g^{i_2 i_6} g^{i_4 i_5} + \\ & g^{i_1 i_6} g^{i_2 i_3} g^{i_4 i_5} + g^{i_1 i_4} g^{i_2 i_5} g^{i_3 i_6} + g^{i_1 i_5} g^{i_2 i_4} g^{i_3 i_6} + g^{i_1 i_4} g^{i_2 i_6} g^{i_3 i_5} + \\ & g^{i_1 i_6} g^{i_2 i_4} g^{i_3 i_5} + g^{i_1 i_5} g^{i_2 i_6} g^{i_3 i_4} + g^{i_1 i_6} g^{i_2 i_5} g^{i_3 i_4}). \end{aligned} \quad (40)$$

Substituting these into (38) and using the fact that the cumulants $\kappa_{i_1 \dots i_r}$ are invariant under all permutations of i_1, \dots, i_r , by (15), then completes the proof.

A.3 Proof of Lemma 3

Let Θ be the natural parameter space, so $\log p_1(x, \theta) = x \cdot \theta - \psi(\theta)$, and let $\partial_i = \partial/\partial\theta_i$. Then

$$g_{ij} = \partial_i \partial_j \psi = -\partial_i \partial_j \log p_1, \quad (41)$$

by [13, Theorem 2.2.5]. In particular, this implies that the right-hand side of (41) is constant in x . Then

$$\begin{aligned} \text{AC}_{i_1, \dots, i_r}^{(r)} &= \int \left(\prod_{j=1}^r \partial_{i_j} \log p_1 \right) p_1 d\mu_1(x) \text{ by (28)} \\ &= \int \left(\prod_{j=1}^{r-1} \partial_{i_j} \log p_1 \right) \partial_{i_r} p_1 d\mu_1(x) \\ &= \partial_{i_r} \int \left(\prod_{j=1}^{r-1} \partial_{i_j} \log p_1 \right) p_1 d\mu_1(x) - \int \partial_{i_r} \left(\prod_{j=1}^{r-1} \partial_{i_j} \log p_1 \right) p_1 d\mu_1(x) \\ &= \partial_{i_r} \text{AC}_{i_1, \dots, i_{r-1}}^{(r-1)} - \int \sum_{k=1}^{r-1} \left((\partial_{i_r} \partial_{i_k} \log p_1) \prod_{j \neq k} \partial_{i_j} \log p_1 \right) p_1 d\mu_1(x) \\ &= \partial_{i_r} \text{AC}_{i_1, \dots, i_{r-1}}^{(r-1)} + \sum_{k=1}^{r-1} g_{i_r i_k} \int \left(\prod_{j \neq k} \partial_{i_j} \log p_1 \right) p_1 d\mu_1(x) \text{ by (41)} \\ &= \partial_{i_r} \text{AC}_{i_1, \dots, i_{r-1}}^{(r-1)} + \sum_{k=1}^{r-1} g_{i_r i_k} \text{AC}_{i_1, \dots, \hat{i}_k, \dots, i_{r-1}}^{(r-2)} \text{ by (28)}. \end{aligned}$$

B Edgeworth expansions and Hermite numbers

This section gives a fairly self-contained description of multivariate Edgeworth expansions and the Hermite numbers for random variables with general (rather than identity) variance-covariance matrices. We use the rigorous error analysis of Bhattacharya and Denker [6] and we closely follow the approach of Kass and Vos [13, §4.5], except that we use multi-index notation which is more convenient when describing objects specified by large numbers of indices (such as high-degree partial derivatives of a function). This section is an exposition of known results only, though its slight reformulation of Edgeworth expansions and its explicit formulae for higher order terms are useful for the rest of this paper, and do not appear in standard references. Our main results are also reformulated in standard tensor notation at the end of this section.

Let Y and Y_1, Y_2, Y_3, \dots be IID d -dimensional random variables with mean 0 and variance-covariance matrix Σ . We assume that the distribution of Y belongs to an exponential family of distributions, and therefore has a moment generating function, but we consider the distribution of Y to be fixed for this entire section. We aim to calculate the Edgeworth expansion for the distribution ζ_n of $Z_n \stackrel{\text{def}}{=} (Y_1 + \dots + Y_n)/\sqrt{n}$ as $n \rightarrow \infty$ in terms of the cumulants of Y .

Let M_Y and K_Y be (respectively) the moment and cumulant generating functions of Y , given by $M_Y(t) \stackrel{\text{def}}{=} \mathbb{E}[e^{t \cdot Y}]$ and $K_Y(t) = \log M_Y(t)$ for any $t \in \mathbb{R}^d$, where the dot represents the Euclidean inner product on \mathbb{R}^d . Any d -tuple $\alpha = (\alpha_1, \dots, \alpha_d)$ of non-negative integers is called a multi-index, and we define $|\alpha| \stackrel{\text{def}}{=} \alpha_1 + \dots + \alpha_d$,

$$\partial_t^\alpha = \frac{\partial^{|\alpha|}}{\partial t_1^{\alpha_1} \dots \partial t_d^{\alpha_d}},$$

$\alpha! = \alpha_1! \dots \alpha_d!$ and $t^\alpha = t_1^{\alpha_1} \dots t_d^{\alpha_d}$ for any $t \in \mathbb{R}^d$. Then the α^{th} cumulant κ_α of Y is defined to be

$$\kappa_\alpha = (\partial_t^\alpha K_Y)|_{t=0}. \quad (42)$$

If e_1, \dots, e_d is the standard basis for \mathbb{R}^d then each e_i is a multi-index, and a direct calculation from the definition shows that $\kappa_{e_i} = \mathbb{E}[Y^{(i)}]$ where $Y^{(i)}$ is the i^{th} component of $Y = (Y^{(1)}, \dots, Y^{(d)})$. We also note for future reference that $\kappa_{e_i+e_j} = \Sigma_{ij}$ (directly from the definition), so the degree-2 cumulants of Y essentially give the covariance matrix of Y . Then by (42), the Taylor series expansion for K_Y about 0 can be written as

$$K_Y(t) = \sum_{|\alpha| \geq 2} \kappa_\alpha t^\alpha / \alpha! \quad (43)$$

where the sum is over all multi-indices α with $|\alpha| \geq 2$, since $\kappa_\alpha = 0$ when $|\alpha|$ is 0 or 1, because $K_Y(0) = 0$ always and $\mathbb{E}[Y] = 0$ by assumption.

Let $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ be the PDF for the d -dimensional normal distribution $N_d(0, \Sigma)$ with the same mean and covariance matrix as Y , so

$$\phi(y) = (2\pi)^{-d/2} (\det \Sigma)^{-1/2} \exp(-y^T \Sigma^{-1} y / 2) \quad (44)$$

for any $y \in \mathbb{R}^d$, where y^T is the transpose of y (here thinking of y as a column matrix). Then define the Hermite polynomial $h_\alpha(y)$ corresponding to multi-index α by

$$h_\alpha(y) \phi(y) = (-1)^{|\alpha|} \partial_y^\alpha \phi(y) \quad (45)$$

where ∂_y^α is a partial differential operator with respect to y which is defined similarly to ∂_t^α . Note that many authors take Σ to be the identity when defining the Hermite polynomials, but Σ will be general throughout this section.

Let \mathcal{A} be a class of Borel subsets of \mathbb{R}^d which satisfies the condition in [6, Corollary 1.4], e.g. \mathcal{A} could consist of all convex Borel sets [6, Remark 1.4.1]. Let λ be the Lebesgue measure on \mathbb{R}^d .

Theorem 4. *Suppose the distribution of Y satisfies Cramér's condition (7). Then for any integral of half-integral $s \geq 0$, there exist polynomials $S_{i/2}(y)$ in y for each $i = 0, \dots, 2s$ so that*

$$\int_B d\zeta_n(y) = \int_B \phi(y) \left[\sum_{i=0}^{2s} S_{i/2}(y) n^{-i/2} \right] d\lambda(y) + o(n^{-s})$$

uniformly in all Borel sets $B \in \mathcal{A}$. In particular, $S_0(y) = 1$,

$$S_1(y) = \sum_{|\alpha|=4} \frac{\kappa_\alpha h_\alpha(y)}{\alpha!} + \sum_{|\alpha|=|\beta|=3} \frac{\kappa_\alpha \kappa_\beta h_{\alpha+\beta}(y)}{2 \alpha! \beta!}$$

and

$$\begin{aligned} S_2(y) = & \sum_{|\alpha|=6} \frac{\kappa_\alpha h_\alpha(y)}{\alpha!} + \sum_{\substack{|\alpha|=3 \\ |\beta|=5}} \frac{\kappa_\alpha \kappa_\beta h_{\alpha+\beta}(y)}{\alpha! \beta!} + \sum_{|\alpha|=|\beta|=4} \frac{\kappa_\alpha \kappa_\beta h_{\alpha+\beta}(y)}{2 \alpha! \beta!} \\ & + \sum_{\substack{|\alpha|=|\beta|=3 \\ |\gamma|=4}} \frac{\kappa_\alpha \kappa_\beta \kappa_\gamma h_{\alpha+\beta+\gamma}(y)}{2 \alpha! \beta! \gamma!} + \sum_{\substack{|\alpha|=|\beta|=3 \\ |\gamma|=|\delta|=3}} \frac{\kappa_\alpha \kappa_\beta \kappa_\gamma \kappa_\delta h_{\alpha+\beta+\gamma+\delta}(y)}{24 \alpha! \beta! \gamma! \delta!}. \end{aligned}$$

Also, if i is odd then $S_{i/2}(y)$ is a linear combination of terms of the form $h_\alpha(y)$ where $|\alpha|$ is odd (which implies $S_{i/2}(y) = 0$ in our main case of interest, $y = 0$, as we will see in Theorem 5, below).

Proof. As mentioned above, this proof closely follows the approach of [13, §4.5] and relies on the rigorous error analysis of [6].

We first note that

$$\begin{aligned}
K_{Z_n}(t) &= nK_Y(t/\sqrt{n}) \text{ since } Y_1, \dots, Y_n \text{ are IID} \\
&= \sum_{|\alpha| \geq 2} \kappa_\alpha t^\alpha n^{1-|\alpha|/2} / \alpha! \text{ by (43)} \\
&= \frac{1}{2} t^T \Sigma t + \sum_{|\alpha| \geq 3} \kappa_\alpha t^\alpha n^{1-|\alpha|/2} / \alpha! \text{ since } \kappa_{e_i+e_j} = \Sigma_{ij} \\
&= \frac{1}{2} t^T \Sigma t + \sum_{r=3}^{\infty} T_r n^{1-r/2}
\end{aligned}$$

where $T_r \stackrel{\text{def}}{=} \sum_{|\alpha|=r} \kappa_\alpha t^\alpha / \alpha!$ is the sum of all terms of degree r in t . So

$$\begin{aligned}
M_{Z_n}(t) &= e^{t^T \Sigma t / 2} \exp \left(\sum_{r=3}^{\infty} T_r n^{1-r/2} \right) \\
&= e^{t^T \Sigma t / 2} \left(1 + \left[\sum_{r=3}^{\infty} T_r n^{1-r/2} \right] + \frac{1}{2!} \left[\sum_{r=3}^{\infty} T_r n^{1-r/2} \right]^2 + \frac{1}{3!} \left[\sum_{r=3}^{\infty} T_r n^{1-r/2} \right]^3 + \dots \right) \\
&= e^{t^T \Sigma t / 2} \left(1 + n^{-1/2} [T_3] + n^{-1} [T_4 + T_3^2/2] + n^{-3/2} [T_5 + T_3 T_4 + T_3^3/6] \right. \\
&\quad + n^{-2} [T_6 + T_3 T_5 + T_4^2/2 + T_3^2 T_4/2 + T_3^4/24] \\
&\quad + n^{-5/2} [T_7 + T_3 T_6 + T_4 T_5 + T_3^2 T_5/2 + T_3 T_4^2/2 + T_3^3 T_4/6 + T_3^5/120] \\
&\quad \left. + O(n^{-3}) \right) \tag{46}
\end{aligned}$$

where the second step uses the Taylor series expansion for the exponential function about 0 (note that the argument is $O(1/\sqrt{n})$ so any desired accuracy can be obtained by truncating this series appropriately) and the third step simply collects together terms with the same power of n (which is easy to do using a program which performs symbolic formula manipulation, such as Maxima [18]).

Now, $M_{Z_n}(t)$ is the integral transform

$$M_{Z_n}(t) = \mathbb{E}[e^{t \cdot Z_n}] = \int_{\mathbb{R}^d} e^{t \cdot y} d\zeta_n(y)$$

of the distribution ζ_n which are trying to approximate, so we need to invert this transform. But (46) approximates $M_{Z_n}(t)$ by a linear combination of terms which are each of the form $t^\alpha e^{t^T \Sigma t / 2}$, which we will shortly show is the transform of $h_\alpha(y) \phi(y)$. Using this property of the Hermite polynomials to invert (46) term-by-term then gives the Edgeworth expansion for ζ_n . For example, the n^{-1} term in (46) is

$$n^{-1} e^{t^T \Sigma t / 2} [T_4 + T_3^2/2] = n^{-1} e^{t^T \Sigma t / 2} \left[\sum_{|\alpha|=4} \frac{\kappa_\alpha t^\alpha}{\alpha!} + \sum_{|\alpha|=|\beta|=3} \frac{\kappa_\alpha \kappa_\beta t^{\alpha+\beta}}{2 \alpha! \beta!} \right]$$

whose inverse transform is therefore $n^{-1} \phi(y) S_1(y)$, where $S_1(y)$ is given in the statement of the theorem. The rigorous error analysis of [6, Corollary 1.4] then proves the theorem.

So we finish by showing that the transform of $h_\alpha(y)\phi(y)$ is $t^\alpha e^{t^T \Sigma t/2}$:

$$\begin{aligned}
& \int_{\mathbb{R}^d} e^{t \cdot y} h_\alpha(y) \phi(y) d\lambda(y) \\
&= \int_{\mathbb{R}^d} e^{t \cdot y} (-1)^{|\alpha|} (\partial_y^\alpha \phi) d\lambda(y) \text{ by the definition (45) of } h_\alpha(y) \\
&= \int_{\mathbb{R}^d} (\partial_y^\alpha e^{t \cdot y}) \phi(y) d\lambda(y) \text{ by integration by parts} \\
&= t^\alpha \int_{\mathbb{R}^d} e^{t \cdot y} \phi(y) d\lambda(y) \\
&= t^\alpha \int_{\mathbb{R}^d} (2\pi)^{-d/2} (\det \Sigma)^{-1/2} \exp(t \cdot y - y^T \Sigma^{-1} y/2) d\lambda(y) \text{ by (44)} \\
&= t^\alpha \exp(t^T \Sigma t/2) \int_{\mathbb{R}^d} (2\pi)^{-d/2} (\det \Sigma)^{-1/2} \exp(-(y - \Sigma t)^T \Sigma^{-1} (y - \Sigma t)/2) d\lambda(y) \\
&= t^\alpha e^{t^T \Sigma t/2},
\end{aligned}$$

where the last step uses the fact that the integrand is the PDF for the normal distribution $N_d(\Sigma t, \Sigma)$ whose integral is 1. \square

When applying Theorem 4 in the rest of this paper, our main case of interest will be when $y = 0$. We therefore now calculate the numbers $h_\alpha(0)$, which are known as the Hermite numbers.

Theorem 5. *The Hermite number $h_\alpha(0)$ is 0 if $|\alpha|$ is odd and if $|\alpha| = 2k$ is even then*

$$h_\alpha(0) = \frac{(-1)^k}{k!} \partial_z^\alpha \left[(z^T \Sigma^{-1} z/2)^k \right] \quad (47)$$

$$= \frac{(-1)^k \alpha!}{k!} \sum_{\substack{|\beta^1|=\dots=|\beta^k|=2 \\ \alpha=\beta^1+\dots+\beta^k}} \frac{\tilde{\kappa}_{\beta^1} \dots \tilde{\kappa}_{\beta^k}}{\beta^1! \dots \beta^k!}, \quad (48)$$

where the sum is over all ordered k -tuples β^1, \dots, β^k of degree-2 multi-indices whose sum is α , and where $\tilde{\kappa}_{\beta^j} = (\Sigma^{-1})_{ab}$ if $\beta^j = e_a + e_b$.

Proof. Recall that $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ is the PDF of $N_d(0, \Sigma)$ given by (44) and that the Hermite polynomial $h_\alpha(y)$ is defined by (45), i.e., $h_\alpha \phi = (-1)^{|\alpha|} (\partial_y^\alpha \phi)$. Using (45), we see that our version of the Hermite polynomials satisfy the following recurrence relation, where we recall that e_1, \dots, e_d is the standard basis for \mathbb{R}^d :

$$h_{\alpha+e_i} = \phi^{-1} (-1)^{|\alpha+e_i|} (\partial_y^{\alpha+e_i} \phi) = -\phi^{-1} \frac{\partial}{\partial y_i} (h_\alpha \phi) = -\frac{\partial h_\alpha}{\partial y_i} + h_\alpha h_{e_i}, \quad (49)$$

where $\phi^{-1} = 1/\phi$. It is easy to show from (44) and (45) that $h_0(y) = 1$ and $h_{e_i}(y) = e_i^T \Sigma^{-1} y$, so the Hermite polynomials can be calculated explicitly from (49), though we will do not pursue this here.

Now, let $z \in \mathbb{R}^d$ be fixed, and define $\langle z, z \rangle = z^T \Sigma^{-1} z$ to be the square of the Mahalanobis norm. Then for $s \in \mathbb{R}$ and any non-negative integer r ,

$$\begin{aligned}
(-1)^r \frac{d^r}{ds^r} [\phi(sz)] &= (2\pi)^{-d/2} (\det \Sigma)^{-1/2} (-1)^r \frac{d^r}{ds^r} [\exp(-s^2 \langle z, z \rangle/2)] \text{ by (44)} \\
&= \langle z, z \rangle^{r/2} \phi(sz) \tilde{h}_r \left(s \sqrt{\langle z, z \rangle} \right)
\end{aligned} \quad (50)$$

where \tilde{h}_r is the r^{th} 1-dimensional, unit variance Hermite polynomial, defined by

$$\tilde{h}_r(w) \exp(-w^2/2) = (-1)^r \frac{d^r}{dw^r} \exp(-w^2/2).$$

On the other hand, we will shortly use the chain rule to show that

$$(-1)^r \frac{d^r}{ds^r} [\phi(sz)] = \phi(sz) \sum_{|\alpha|=r} n_\alpha h_\alpha(sz) z^\alpha \quad (51)$$

where $n_\alpha = |\alpha|!/\alpha!$. Combining (50) and (51) and evaluating at $s = 0$ will therefore give

$$\langle z, z \rangle^{r/2} \tilde{h}_r(0) = \sum_{|\alpha|=r} n_\alpha h_\alpha(0) z^\alpha. \quad (52)$$

But it is well known that $\tilde{h}_r(0)$ is 0 if r is odd and $(-1)^{r/2}(r-1)!!$ if r is even, where $(r-1)!! = 1 \times 3 \times 5 \times \cdots \times (r-1)$ is the double factorial. So differentiating both sides of (52) by ∂_z^α gives $h_\alpha(0) = 0$ if $|\alpha|$ is odd and (47) if $|\alpha|$ is even. Then (48) follows from (47) and

$$\frac{z^T \Sigma^{-1} z}{2} = \sum_{|\beta|=2} \frac{\tilde{\kappa}_\beta z^\beta}{\beta!}.$$

So we finish by proving (51) by induction. This is clearly true when $r = 0$ since $h_0(y) = 1$, so we now assume (51) is true for $r = k$ and prove it for $r = k + 1$.

$$\begin{aligned} (-1)^{k+1} \frac{d^{k+1}}{ds^{k+1}} [\phi(sz)] &= -\frac{d}{ds} \left[\phi(sz) \sum_{|\alpha|=k} n_\alpha h_\alpha(sz) z^\alpha \right] \text{ by the induction hypothesis} \\ &= -\sum_{|\alpha|=k} \sum_{i=1}^d z_i \frac{\partial}{\partial y_i} \bigg|_{y=sz} [\phi(y) h_\alpha(y)] n_\alpha z^\alpha \text{ by the chain rule} \\ &= \phi(sz) \sum_{|\alpha|=k} \sum_{i=1}^d n_\alpha z^{\alpha+e_i} \left[h_{e_i}(sz) h_\alpha(sz) - \frac{\partial h_\alpha}{\partial y_i} \bigg|_{y=sz} \right] \text{ by (45)} \\ &= \phi(sz) \sum_{|\alpha|=k} \sum_{i=1}^d n_\alpha z^{\alpha+e_i} h_{\alpha+e_i}(sz) \text{ by (49)} \\ &= \phi(sz) \sum_{|\beta|=k+1} n_\beta h_\beta(sz) z^\beta \end{aligned}$$

since, for any β with $|\beta| = k + 1$, we have

$$\sum_{i: \beta_i \geq 1} n_{\beta-e_i} = \sum_{i: \beta_i \geq 1} \frac{(|\beta| - 1)!}{\beta_1! \cdots (\beta_i - 1)! \cdots \beta_d!} = \frac{(|\beta| - 1)!}{\beta!} \sum_{i: \beta_i \geq 1} \beta_i = \frac{|\beta|!}{\beta!} = n_\beta.$$

□

Multi-index notation is generally more concise than tensor notation for expressions involving multi-indices α where $|\alpha|$ is large, but not when $|\alpha|$ is small. Also, many people are uncomfortable with multi-index notation, despite its advantages, so we finish this section by converting some of the expressions above into tensor notation. The key identity for this conversion is

$$\sum_{|\alpha|=r} \frac{f(\alpha)}{\alpha!} = \frac{1}{r!} \sum_{i_1, \dots, i_r} f(e_{i_1} + \cdots + e_{i_r}), \quad (53)$$

where f is any function of degree r multi-indices. This formula follows from

$$\sum_{i_1, \dots, i_r} f(e_{i_1} + \cdots + e_{i_r}) = \sum_{|\alpha|=r} \sum_{\substack{i_1, \dots, i_r \\ \alpha = e_{i_1} + \cdots + e_{i_r}}} f(\alpha) = \sum_{|\alpha|=r} f(\alpha) \sum_{\substack{i_1, \dots, i_r \\ \alpha = e_{i_1} + \cdots + e_{i_r}}} 1 = \sum_{|\alpha|=r} \frac{f(\alpha) r!}{\alpha!},$$

where the last step uses the fact that $r!/α!$ is the multinomial coefficient which counts the number of ways of putting r distinguishable objects into d boxes so that $α_i$ objects go into box i .

For any indices $i_1, \dots, i_r \in \{1, \dots, d\}$, let $\kappa_{i_1 \dots i_r}$ and $h_{i_1 \dots i_r}$ be alternative notation for κ_α and $h_\alpha(0)$, respectively, where $\alpha = e_{i_1} + \dots + e_{i_r}$. Then (53) implies that $S_1(y)$ and $S_2(y)$ in Theorem 4 are given at $y = 0$ by

$$S_1(0) = \frac{1}{24} \sum_{i_1, \dots, i_4} \kappa_{i_1 \dots i_4} h_{i_1 \dots i_4} + \frac{1}{72} \sum_{i_1, \dots, i_6} \kappa_{i_1 i_2 i_3} \kappa_{i_4 i_5 i_6} h_{i_1 \dots i_6} \quad (54)$$

and

$$\begin{aligned} S_2(0) = & \frac{1}{720} \sum_{i_1, \dots, i_6} \kappa_{i_1 \dots i_6} h_{i_1 \dots i_6} + \frac{1}{720} \sum_{i_1, \dots, i_8} \kappa_{i_1 i_2 i_3} \kappa_{i_4 \dots i_8} h_{i_1 \dots i_8} + \\ & \frac{1}{1152} \sum_{i_1, \dots, i_8} \kappa_{i_1 \dots i_4} \kappa_{i_5 \dots i_8} h_{i_1 \dots i_8} + \frac{1}{1728} \sum_{i_1, \dots, i_{10}} \kappa_{i_1 i_2 i_3} \kappa_{i_4 i_5 i_6} \kappa_{i_7 \dots i_{10}} h_{i_1 \dots i_{10}} + \\ & \frac{1}{31104} \sum_{i_1, \dots, i_{12}} \kappa_{i_1 i_2 i_3} \kappa_{i_4 i_5 i_6} \kappa_{i_7 i_8 i_9} \kappa_{i_{10} i_{11} i_{12}} h_{i_1 \dots i_{12}}. \end{aligned} \quad (55)$$

Also, by Theorem 5 and (53), we have

$$h_{i_1 \dots i_4} = g^{i_1 i_2} g^{i_3 i_4} + g^{i_1 i_3} g^{i_2 i_4} + g^{i_1 i_4} g^{i_2 i_3} \quad (56)$$

and

$$\begin{aligned} h_{i_1 \dots i_6} = & -(g^{i_1 i_2} g^{i_3 i_4} g^{i_5 i_6} + g^{i_1 i_3} g^{i_2 i_4} g^{i_5 i_6} + g^{i_1 i_4} g^{i_2 i_3} g^{i_5 i_6} + g^{i_1 i_2} g^{i_3 i_5} g^{i_4 i_6} + \\ & g^{i_1 i_3} g^{i_2 i_5} g^{i_4 i_6} + g^{i_1 i_5} g^{i_2 i_3} g^{i_4 i_6} + g^{i_1 i_2} g^{i_3 i_6} g^{i_4 i_5} + g^{i_1 i_3} g^{i_2 i_6} g^{i_4 i_5} + \\ & g^{i_1 i_6} g^{i_2 i_3} g^{i_4 i_5} + g^{i_1 i_4} g^{i_2 i_5} g^{i_3 i_6} + g^{i_1 i_5} g^{i_2 i_4} g^{i_3 i_6} + g^{i_1 i_4} g^{i_2 i_6} g^{i_3 i_5} + \\ & g^{i_1 i_6} g^{i_2 i_4} g^{i_3 i_5} + g^{i_1 i_5} g^{i_2 i_6} g^{i_3 i_4} + g^{i_1 i_6} g^{i_2 i_5} g^{i_3 i_4}) \end{aligned} \quad (57)$$

where $g^{ij} = (\Sigma^{-1})_{ij}$. These all agree with [13, §4.5] to the extent that our equations and theirs overlap.

References

- [1] S. Amari and H. Nagaoka. *Methods of Information Geometry*, volume 191 of *Translations of mathematical monographs*. American Mathematical Society, 2000.
- [2] O. Barndorff-Nielsen. *Information and exponential families*. John Wiley & Sons, 1978.
- [3] O. Barndorff-Nielsen and D. R. Cox. Edgeworth and saddle-point approximations with statistical applications. *Journal of the Royal Statistical Society Series B (Methodological)*, 41(3):279–312, 1979.
- [4] A. R. Barron and T. M. Cover. Minimum complexity density estimation. *IEEE Transactions on Information Theory*, 37(4):1034–1054, July 1991.
- [5] A. R. Barron, J. Rissanen, and B. Yu. The minimum description length principle in coding and modeling. *IEEE Transactions on Information Theory*, 44(6):2743–2760, October 1998.
- [6] R. Bhattacharya and M. Denker. *Asymptotic Statistics*, volume 14 of *Oberwolfach Seminars*. Birkhäuser Basel, 1990.

- [7] B. S. Clarke and A. R. Barron. Information-theoretic asymptotics of Bayes methods. *IEEE Transactions on Information Theory*, 36(3):453–471, May 1990.
- [8] B. S. Clarke and A. R. Barron. Jeffreys’ prior is asymptotically least favorable under entropy risk. *Journal of Statistical Planning and Inference*, 41(1):37 – 60, 1994. ISSN 0378-3758.
- [9] B. Efron. Defining the curvature of a statistical problem (with applications to second order efficiency). *Ann. Statist.*, 3(6):1189–1242, 11 1975.
- [10] A. Gray. The volume of a small geodesic ball of a Riemannian manifold. *Michigan Math. J.*, 20:329–344, 1973.
- [11] P. Grünwald. A tutorial introduction to the minimum description length principle. In I. J. M. P. Grünwald and M. Pitt, editors, *Advances in Minimum Description Length: Theory and Applications*. MIT Press, 2005.
- [12] P. Hall. *The bootstrap and Edgeworth expansions*. Springer, New York, 1992.
- [13] R. E. Kass and P. W. Vos. *Geometrical Foundations of Asymptotic Inference*. John Wiley & Sons, 1997.
- [14] E. Nelson. Probability theory and Euclidean field theory. In G. Velo and A. Wightman, editors, *Constructive Quantum Field Theory*, volume 25 of *Lecture Notes in Physics*, pages 94–124. Springer Berlin Heidelberg, 1973.
- [15] J. Rissanen. Fisher information and stochastic complexity. *IEEE Transactions on Information Theory*, 42(1):40–47, January 1996.
- [16] J. Rissanen. MDL denoising. *IEEE Transactions on Information Theory*, 46(7): 2537–2543, 2000.
- [17] J. Rissanen. *Information and Complexity in Statistical Modeling*. Information Science and Statistics. Springer, first edition, 2007.
- [18] W. Schelter. *Maxima, a Computer Algebra System. Version 5.31.2*. Macsyma group, 2015. URL <http://maxima.sourceforge.net>.
- [19] Y. M. Shtarkov. Universal sequential coding of single messages. *Probl. Inform. Transm.*, 23(3):3–17, 1987.